

(12) **United States Patent**
Mascaro et al.

(10) **Patent No.:** **US 9,208,794 B1**
(45) **Date of Patent:** **Dec. 8, 2015**

(54) **PROVIDING SOUND MODELS OF AN INPUT SIGNAL USING CONTINUOUS AND/OR LINEAR FITTING**

7,389,230 B1	6/2008	Nelken	704/255
7,664,640 B2	2/2010	Webber	704/243
7,668,711 B2	2/2010	Chong et al.	704/219
8,015,002 B2 *	9/2011	Li et al.	704/226
2004/0066940 A1 *	4/2004	Amir	381/94.2

(71) Applicant: **THE INTELLISIS CORPORATION**,
San Diego, CA (US)

(Continued)

(72) Inventors: **Massimo Mascaro**, San Diego, CA
(US); **David C. Bradley**, La Jolla, CA
(US); **Yao Huang Morin**, San Diego,
CA (US)

FOREIGN PATENT DOCUMENTS

WO	WO 2012/129255	9/2012
WO	WO 2012/134991	10/2012
WO	WO 2012/134993	10/2012

(73) Assignee: **The Intellis Corporation**, San Diego,
CA (US)

OTHER PUBLICATIONS

Kumar et al., "Speaker Recognition Using GMM", *International Journal of Engineering Science and Technology*, vol. 2, No. 6, 2010, retrieved from the Internet: <http://www.ijest.info/docs/IJEST10-02-06-112.pdf>, pp. 2428-2436.

(Continued)

(21) Appl. No.: **13/961,811**

(22) Filed: **Aug. 7, 2013**

Primary Examiner — Jesse Pullias

(74) *Attorney, Agent, or Firm* — Edell, Shapiro & Finnan, LLC

(51) **Int. Cl.**
G10L 21/00 (2013.01)
G10L 25/90 (2013.01)
G10L 21/0208 (2013.01)
G10L 21/003 (2013.01)

(57) **ABSTRACT**

(52) **U.S. Cl.**
CPC **G10L 21/0208** (2013.01); **G10L 21/003**
(2013.01)

Voice enhancement and/or speech features extraction may be performed on noisy audio signals. An input signal may convey audio comprising a speech component superimposed on a noise component. The input signal may be segmented into discrete successive time windows including a first time window spanning a duration greater than a sampling interval of the input signal. A transform may be performed on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows. A first sound model may describe a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal. Linear fits in time of the sound models over individual time windows of the input signal may be obtained. The linear fits may include a first linear fit in time of the first sound model over the first time window.

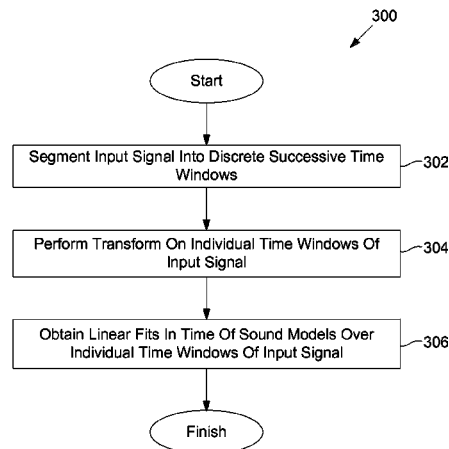
(58) **Field of Classification Search**
USPC 704/200–257, 500–504
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,815,580 A	9/1998	Craven et al.	381/58
5,978,824 A *	11/1999	Ikeda	708/322
6,594,585 B1	7/2003	Gersztenkorn	
7,117,149 B1	10/2006	Zakarauskas	704/233
7,249,015 B2	7/2007	Jiang et al.	704/222

20 Claims, 3 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2004/0128130	A1	7/2004	Rose et al.	704/236
2004/0158462	A1	8/2004	Rutledge et al.	
2004/0167777	A1 *	8/2004	Hetherington et al.	704/226
2004/0176949	A1	9/2004	Wenndt et al.	704/203
2004/0220475	A1	11/2004	Szabo et al.	600/458
2005/0114128	A1	5/2005	Hetherington et al.	704/233
2005/0149321	A1	7/2005	Kabi et al.	
2006/0053003	A1	3/2006	Suzuki et al.	
2006/0100866	A1	5/2006	Alewine et al.	704/226
2006/0100868	A1 *	5/2006	Hetherington et al.	704/226
2006/0136203	A1 *	6/2006	Ichikawa	704/226
2007/0010997	A1	1/2007	Kim	704/208
2008/0082323	A1	4/2008	Bai et al.	704/214
2009/0012638	A1	1/2009	Lou	700/94
2009/0076822	A1	3/2009	Sanjaume	704/268
2010/0260353	A1	10/2010	Ozawa	381/94.3
2010/0299144	A1	11/2010	Barzelay et al.	
2010/0332222	A1	12/2010	Bai et al.	704/214
2011/0016077	A1	1/2011	Vasilache et al.	706/52
2011/0060564	A1	3/2011	Hoge	703/2
2011/0286618	A1	11/2011	Vandali et al.	381/320

2012/0191450	A1	7/2012	Pinson	
2012/0243694	A1	9/2012	Bradley et al.	381/56
2012/0243705	A1	9/2012	Bradley et al.	381/94.4
2012/0243707	A1	9/2012	Bradley et al.	381/98
2013/0255473	A1	10/2013	Abe et al.	

OTHER PUBLICATIONS

Kamath et al, "Independent Component Analysis for Audio Classification", *IEEE 11th Digital Signal Processing Workshop & IEEE Signal Processing Education Workshop*, 2004, retrieved from the Internet: <http://2002.114.89.42/resource/pdf/1412.pdf>, pp. 352-355.

Vargas-Rubio et al., "An Improved Spectrogram Using the Multiangle Centered Discrete Fractional Fourier Transform", *Proceedings of International Conference on Acoustics, Speech, and Signal Processing*, Philadelphia, 2005, retrieved from the internet: <URL: <http://www.ece.unm.edu/faculty/beanthan/PUB/ICASSP-05-JUAN.pdf>>, 4 pages.

U.S. Appl. No. 13/945,731, filed Jul. 18, 2013, 33 pages.

U.S. Appl. No. 13/945,731 Office Action dated Jan. 1, 2015, citing prior art, 12 pages.

* cited by examiner

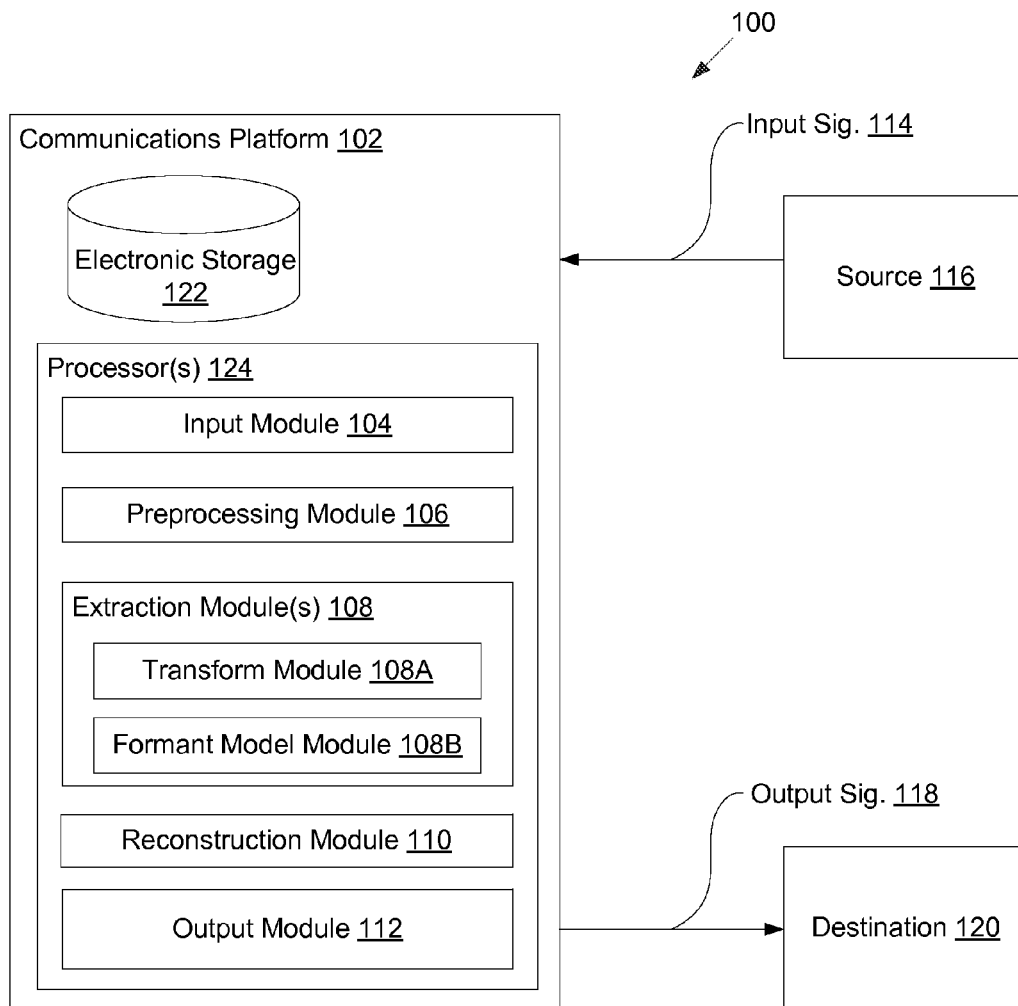


FIG. 1




FIG. 2

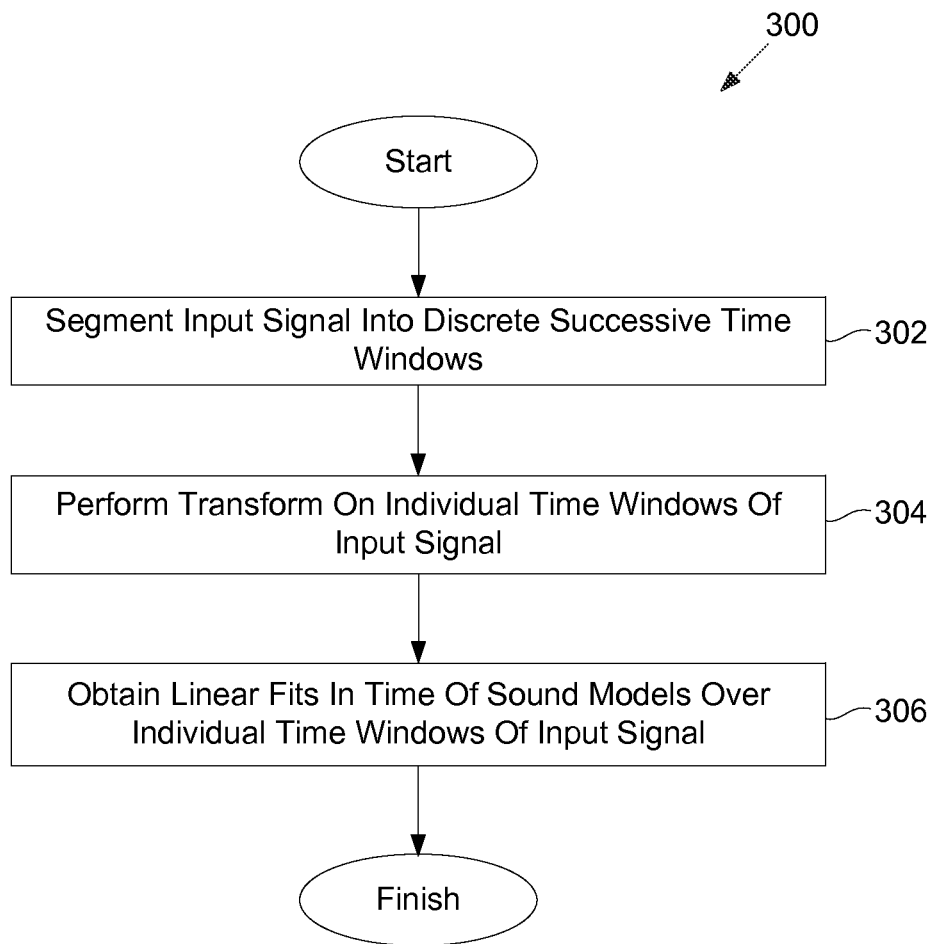


FIG. 3

1

PROVIDING SOUND MODELS OF AN INPUT SIGNAL USING CONTINUOUS AND/OR LINEAR FITTING

FIELD OF THE DISCLOSURE

This disclosure relates to providing sound models of an input signal using continuous and/or linear fitting.

BACKGROUND

Systems configured to identify speech in an audio signal are known. Existing systems, however, typically may rely on an ability to identify phonemes in the signal. A phoneme-based approach may be unreliable at least because phonemes may vary according to context.

SUMMARY

One aspect of the disclosure relates to a system configured to perform voice enhancement and/or speech features extraction on noisy audio signals, in accordance with one or more implementations. Voice enhancement and/or speech features extraction may be performed on noisy audio signals. An input signal may convey audio comprising a speech component superimposed on a noise component. The input signal may be segmented into discrete successive time windows including a first time window spanning a duration greater than a sampling interval of the input signal. A transform may be performed on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows. A first sound model may describe a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal. Linear fits in time of the sound models over individual time windows of the input signal may be obtained. The linear fits may include a first linear fit in time of the first sound model over the first time window.

The communications platform may be configured to execute computer program modules. The computer program modules may include one or more of an input module, one or more extraction modules, a reconstruction module, an output module, and/or other modules.

The input module may be configured to receive an input signal from a source. The input signal may include human speech (or some other wanted signal) and noise. The waveforms associated with the speech and noise may be superimposed in input signal.

Generally speaking, the extraction module(s) may be configured to extract harmonic information from the input signal. The extraction module(s) may include one or more of a transform module, a formant model module, and/or other modules.

The transform module may be configured to perform a transform on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows. A first sound model may describe a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal. Pitch may be the rate of change of phase over time. Chirp may be the rate of change of pitch over time.

The transform module may be configured to obtain linear fits in time of the sound models over individual time windows of the input signal. The linear fits may include a first linear fit in time of the first sound model over the first time window. In some implementations, a linear regression may be used to fit the first sound model over the first time window to obtain the first linear fit. The first model may be assumed to be a super-

2

position of harmonics in the first time window with a linearly varying fundamental frequency. Harmonic amplitudes in the first sound model may be assumed to be piecewise linear in time.

According to some implementations, the transform module may be configured to impose continuity in a pitch estimation of the first sound model. An integral phase of the first sound model may be optimized via a nonlinear regression. The integral phase may be optimized via multiple iterations of the nonlinear regression. A regression to estimate the integral phase may be performed locally, in some implementations. The integral phase may be approximated with a number of time points to reduce the degrees of freedom.

In some implementations, the transform module may be configured to impose continuity in harmonic amplitudes and/or phase estimation. An estimation of harmonic amplitudes and/or phase may be optimized via a nonlinear regression. The estimation of harmonic amplitudes and/or phase may be optimized via multiple iterations of the nonlinear regression. The estimation of harmonic amplitudes and/or phase may be performed locally, in some implementations. The estimation of the harmonic amplitudes and/or phase may be approximated with a number of time points to reduce the degrees of freedom.

The formant model module may be configured to model harmonic amplitudes based on a formant model. Generally speaking, a formant may be described as the spectral resonance peaks of the sound spectrum of the voice. One formant model—the source-filter model—postulates that vocalization in humans occurs via an initial periodic signal produced by the glottis (i.e., the source), which is then modulated by resonances in the vocal and nasal cavities (i.e., the filter).

The reconstruction module may be configured to reconstruct the speech component of the input signal with the noise component of the input signal being suppressed. The reconstruction may be performed once each of the parameters of the formant model has been determined. The reconstruction may be performed by interpolating all the time-dependent parameters and then resynthesizing the waveform of the speech component of the input signal.

The output module may be configured to transmit an output signal to a destination. The output signal may include the reconstructed speech component of the input signal.

These and other features, and characteristics of the present technology, as well as the methods of operation and functions of the related elements of structure and the combination of parts and economies of manufacture, will become more apparent upon consideration of the following description and the appended claims with reference to the accompanying drawings, all of which form a part of this specification, wherein like reference numerals designate corresponding parts in the various figures. It is to be expressly understood, however, that the drawings are for the purpose of illustration and description only and are not intended as a definition of the limits of the invention. As used in the specification and in the claims, the singular form of “a”, “an”, and “the” include plural referents unless the context clearly dictates otherwise.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates a system configured to perform voice enhancement and/or speech features extraction on noisy audio signals, in accordance with one or more implementations.

FIG. 2 illustrates an exemplary spectrogram, in accordance with one or more implementations.

FIG. 3 illustrates a method for performing voice enhancement and/or speech features extraction on noisy audio signals, in accordance with one or more implementations.

DETAILED DESCRIPTION

Voice enhancement and/or speech features extraction may be performed on noisy audio signals. An input signal may convey audio comprising a speech component superimposed on a noise component. The input signal may be segmented into discrete successive time windows including a first time window spanning a duration greater than a sampling interval of the input signal. A transform may be performed on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows. A first sound model may describe a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal. Linear fits in time of the sound models over individual time windows of the input signal may be obtained. The linear fits may include a first linear fit in time of the first sound model over the first time window.

FIG. 1 illustrates a system **100** configured to perform voice enhancement and/or speech features extraction on noisy audio signals, in accordance with one or more implementations. Voice enhancement may be also referred to as denoising or voice cleaning. As depicted in FIG. 1, system **100** may include a communications platform **102** and/or other components. Generally speaking, a noisy audio signal containing speech may be received by communications platform **102**. The communications platform **102** may extract harmonic information from the noisy audio signal. The harmonic information may be used to reconstruct speech contained in the noisy audio signal. By way of non-limiting example, communications platform **102** may include a mobile communications device such as a smart phone, according to some implementations. Other types of communications platforms are contemplated by the disclosure, as described further herein.

The communications platform **102** may be configured to execute computer program modules. The computer program modules may include one or more of an input module **104**, a preprocessing module **106**, one or more extraction modules **108**, a reconstruction module **110**, an output module **112**, and/or other modules.

The input module **104** may be configured to receive an input signal **114** from a source **116**. The input signal **114** may include human speech (or some other wanted signal) and noise. The waveforms associated with the speech and noise may be superimposed in input signal **114**. The input signal **114** may include a single channel (i.e., mono), two channels (i.e., stereo), and/or multiple channels. The input signal **114** may be digitized.

Speech is the vocal form of human communication. Speech is based upon the syntactic combination of lexicals and names that are drawn from very large vocabularies (usually in the range of about 10,000 different words). Each spoken word is created out of the phonetic combination of a limited set of vowel and consonant speech sound units. Normal speech is produced with pulmonary pressure provided by the lungs which creates phonation in the glottis in the larynx that is then modified by the vocal tract into different vowels and consonants. Various differences among vocabularies, syntax that structures individual vocabularies, sets of speech sound units associated with individual vocabularies, and/or other differences create the existence of many thousands of different types of mutually unintelligible human languages.

The noise included in input signal **114** may include any sound information other than a primary speaker's voice. The noise included in input signal **114** may include structured noise and/or unstructured noise. A classic example of structured noise may be a background scene where there are multiple voices, such as a café or a car environment. Unstructured noise may be described as noise with a broad spectral density distribution. Examples of unstructured noise may include white noise, pink noise, and/or other unstructured noise. White noise is a random signal with a flat power spectral density. Pink noise is a signal with a power spectral density that is inversely proportional to the frequency.

An audio signal, such as input signal **114**, may be visualized by way of a spectrogram. A spectrogram is a time-varying spectral representation that shows how the spectral density of a signal varies with time. Spectrograms may be referred to as spectral waterfalls, sonograms, voiceprints, and/or voicegrams. Spectrograms may be used to identify phonetic sounds. FIG. 2 illustrates an exemplary spectrogram **200**, in accordance with one or more implementations. In spectrogram **200**, the horizontal axis represents time (t) and the vertical axis represents frequency (f). A third dimension indicating the amplitude of a particular frequency at a particular time emerges out of the page. A trace of an amplitude peak as a function of time may delineate a harmonic in a signal visualized by a spectrogram (e.g., harmonic **202** in spectrogram **200**). In some implementations, amplitude may be represented by the intensity or color of individual points in a spectrogram. In some implementations, a spectrogram may be represented by a 3-dimensional surface plot. The frequency and/or amplitude axes may be either linear or logarithmic, according to various implementations. An audio signal may be represented with a logarithmic amplitude axis (e.g., in decibels, or dB), and a linear frequency axis to emphasize harmonic relationships or a logarithmic frequency axis to emphasize musical, tonal relationships.

Referring again to FIG. 1, source **116** may include a microphone (i.e., an acoustic-to-electric transducer), a remote device, and/or other source of input signal **114**. By way of non-limiting illustration, where communications platform **102** is a mobile communications device, a microphone integrated in the mobile communications device may provide input signal **114** by converting sound from a human speaker and/or sound from an environment of communications platform **102** into an electrical signal. As another illustration, input signal **114** may be provided to communications platform **102** from a remote device. The remote device may have its own microphone that converts sound from a human speaker and/or sound from an environment of the remote device. The remote device may be the same as or similar to communications platforms described herein.

The preprocessing module **106** may be configured to segment input signal **114** into discrete successive time windows. According to some implementations, a given time window may have a duration in the range of 30-60 milliseconds. In some implementations, a given time window may have a duration that is shorter than 30 milliseconds or longer than 60 milliseconds. The individual time windows of segmented input signal **114** may have equal durations. In some implementations, the duration of individual time windows of segmented input signal **114** may be different. For example, the duration of a given time window of segmented input signal **114** may be based on the amount and/or complexity of audio information contained in the given time window such that the duration increases responsive to a lack of audio information or a presence of stable audio information (e.g., a constant tone).

5

Generally speaking, extraction module(s) **108** may be configured to extract harmonic information from input signal **114**. The extraction module(s) **108** may include one or more of a transform module **108A**, a formant model module **108B**, and/or other modules.

The transform module **108A** may be configured to obtain a sound model over individual time windows of input signal **114**. In some implementations, transform module **108A** may be configured to obtain a linear fit in time of a sound model over individual time windows of input signal **114**. A sound model may be described as a mathematical representation of harmonics in an audio signal. A harmonic may be described as a component frequency of the audio signal that is an integer multiple of the fundamental frequency (i.e., the lowest frequency of a periodic waveform or pseudo-periodic waveform). That is, if the fundamental frequency is f , then harmonics have frequencies $2f$, $3f$, $4f$, etc.

The transform module **108A** may be configured to model input signal **114** as a superposition of harmonics that all share a common pitch and chirp. Such a model may be expressed as:

$$m(t) = 2\Re \left(\sum_{h=1}^{N_h} A_h e^{j2\pi h(\phi t + \frac{\chi}{2} t^2)} \right), \quad \text{EQN. 1}$$

where ϕ is the base pitch and χ is the fractional chirp rate

$$\left(\chi = \frac{c}{\phi}, \text{ where } c \text{ is the actual chirp} \right),$$

where c is the actual chirp), both assumed to be constant. Pitch is defined as the rate of change of phase over time. Chirp is defined as the rate of change of pitch (i.e., the second time derivative of phase). The model of input signal **114** may be assumed as a superposition of N_h harmonics with a linearly varying fundamental frequency. A_h is a complex coefficient weighting all the different harmonics. Being complex, A_h carries information about both the amplitude and about the initial phase for each harmonic.

The model of input signal **114** as a function of A_h may be linear, according to some implementations. In such implementations, linear regression may be used to fit the model, such as follows:

$$\sum_{h=1}^{N_h} A_h e^{j2\pi h(\phi t + \frac{\chi}{2} t^2)} = M(\phi, \chi, t) \bar{A} \quad \text{EQN. 2}$$

with, discretizing time as $(t_1, t_2, \dots, t_{N_t})$:

$$M(\phi, \chi) =$$

$$\begin{bmatrix} e^{j2\pi(\phi t_1 + \frac{\chi}{2} t_1^2)} & e^{j2\pi 2(\phi t_1 + \frac{\chi}{2} t_1^2)} & \dots & e^{j2\pi N_h(\phi t_1 + \frac{\chi}{2} t_1^2)} \\ e^{j2\pi(\phi t_2 + \frac{\chi}{2} t_2^2)} & e^{j2\pi 2(\phi t_2 + \frac{\chi}{2} t_2^2)} & \dots & e^{j2\pi N_h(\phi t_2 + \frac{\chi}{2} t_2^2)} \\ \vdots & \vdots & \ddots & \vdots \\ e^{j2\pi(\phi t_{N_t} + \frac{\chi}{2} t_{N_t}^2)} & e^{j2\pi 2(\phi t_{N_t} + \frac{\chi}{2} t_{N_t}^2)} & \dots & e^{j2\pi N_h(\phi t_{N_t} + \frac{\chi}{2} t_{N_t}^2)} \end{bmatrix}$$

$$\bar{A} = \begin{pmatrix} A_1 \\ \vdots \\ A_{N_h} \end{pmatrix}$$

6

The best value for \bar{A} may be solved via standard linear regression in discrete time, as follows:

$$\bar{A} = M(\phi, \chi) \backslash s, \quad \text{EQN. 3}$$

where the symbol \backslash represents matrix left division (e.g., linear regression).

Due to input signal **114** being real, the fitted coefficients may be doubled with their complex conjugates as:

$$m(t) = (M(\phi, \chi) \ M^*(\phi, \chi)) \begin{pmatrix} \bar{A} \\ \bar{A}^* \end{pmatrix}. \quad \text{EQN. 4}$$

The optimal values of ϕ, χ may not be determinable via linear regression. A nonlinear optimization step may be performed to determine the optimal values of ϕ, χ . Such a nonlinear optimization may include using the residual sum of squares as the optimization metric:

$$[\hat{\phi}, \hat{\chi}] = \underset{\phi, \chi}{\operatorname{argmin}} \left[\sum_t (s(t), \phi, \chi, \bar{A})^2 \right]_{\bar{A} = M(\phi, \chi) \backslash s}, \quad \text{EQN. 5}$$

where the minimization is performed on ϕ, χ at the value of \bar{A} given by the linear regression for each value of the parameters being optimized.

The transform module **108A** may be configured to impose continuity to different fits over time. That is, both continuity in the pitch estimation and continuity in the coefficients estimation may be imposed to extend the model set forth in EQN. 1. If the pitch becomes a continuous function of time (i.e., $\phi = \phi(t)$), then the chirp may be not needed because the fractional chirp may be determined by the derivative of $\phi(t)$ as

$$\chi(t) = \frac{1}{\phi(t)} \frac{d\phi(t)}{dt}.$$

According to some implementations, the model set forth by EQN. 1 may be extended to accommodate a more general time dependent pitch as follows:

$$m(t) = \Re \left(\sum_{h=1}^{N_h} A_h(t) e^{j2\pi h \int_0^t \phi(\tau) d\tau} \right) = \Re \left(\sum_{h=1}^{N_h} A_h(t) e^{j h \Phi(t)} \right), \quad \text{EQN. 6}$$

where $\Phi(t) = 2\pi \int_0^t \phi(\tau) d\tau$ is integral phase.

According to model set forth in EQN. 6, the harmonic amplitudes $A_h(t)$ are time dependent. The harmonic amplitudes may be assumed to be piecewise linear and/or continuous in time such that linear regression may be invoked to obtain $A_h(t)$ for a given integral phase $\Phi(t)$:

$$A_h(t) = A_h(0) + \sum_i \Delta A_h^i \sigma \left(\frac{t - t^{i-1}}{t - t^{i-1}} \right), \quad \text{EQN. 7}$$

where

$$\sigma(t) = \begin{cases} 0 & \text{for } t < 0 \\ t & \text{for } 0 \leq t \leq 1 \\ 1 & \text{for } t > 1 \end{cases}$$

and ΔA_h^i are time-dependent harmonic coefficients. The time-dependent harmonic coefficients ΔA_h^i represent the variation on the complex amplitudes at times t^i .

EQN. 7 may be substituted into EQN. 6 to obtain a linear function of the time-dependent harmonic coefficients ΔA_h^i . The time-dependent harmonic coefficients ΔA_h^i may be solved using standard linear regression for a given integral phase $\Phi(t)$. Actual amplitudes may be reconstructed by

$$A_h^i = A_h^0 + \sum_1^i \Delta A_h^i.$$

The linear regression may be determined efficiently due to the fact that the correlation matrix of the model associated with EQN. 6 and EQN. 7 has a block Toeplitz structure, in accordance with some implementations.

A given integral phase $\Phi(t)$ may be optimized via nonlinear regression. Such a nonlinear regression may be performed using a metric similar to EQN. 5. In order to reduce the degrees of freedom, $\Phi(t)$ may be approximated with a number of time points across which to interpolate by $\Phi(t) = \text{interp}(\Phi^1 = \Phi(t^1), \Phi^2 = \Phi(t^2), \dots, \Phi^{N_t} = \Phi(t^{N_t}))$. In some implementations, the interpolation function may be cubic. The nonlinear optimization of the integral pitch may be:

$$[\Phi^1, \Phi^{N_t}, \dots, \Phi^{N_t}] = \underset{\Phi^1, \Phi^2, \dots, \Phi^{N_t}}{\text{argmin}} \left[\sum_t (s(t) - m(t, \Phi(t), \overline{A_h^i}))^2 \right]_{\substack{\overline{A_h^i} = M(\Phi(t)) \cdot \overline{u(t)} \\ \Phi(t) = \text{interp}(\Phi^1, \Phi^2, \dots, \Phi^{N_t})}} \quad \text{EQN. 8}$$

The different Φ^i may be optimized one at a time with multiple iterations across them. Because each Φ^i affects the integral phase only around t^i , the optimization may be performed locally, according to some implementations.

The formant model module 108B may be configured to model harmonic amplitudes based on a formant model. Generally speaking, a formant may be described as the spectral resonance peaks of the sound spectrum of the voice. One formant model—the source-filter model—postulates that vocalization in humans occurs via an initial periodic signal produced by the glottis (i.e., the source), which is then modulated by resonances in the vocal and nasal cavities (i.e., the filter). In some implementations, the harmonic amplitudes may be modeled according to the source-filter model as:

$$A_h(t) = A(t)G(g(t), \omega(t)) \left[\prod_{r=1}^{N_f} F(f_r(t), \omega(t)) \right] R(\omega(t)) \Big|_{\omega(t)=\hat{\phi}(t)h} \quad \text{EQN. 9}$$

where $A(t)$ is a global amplitude scale common to all the harmonics, but time dependent. G characterizes the source as a function of glottal parameters $g(t)$. Glottal parameters $g(t)$ may be a vector of time dependent parameters. In some implementations, G may be the Fourier transform of the glottal pulse. F describes a resonance (e.g., a formant). The various

cavities in a vocal tract may generate a number of resonances F that act in series. Individual formants may be characterized by a complex parameter $f_r(t)$. R represents a parameter-independent filter that accounts for the air impedance.

In some implementations, the individual formant resonances may be approximated as single pole transfer functions:

$$F(f(t), \omega(t)) = \frac{f(t)f(t)^*}{(j\omega(t) - f(t))(j\omega(t) - f(t)^*)} \quad \text{EQN. 10}$$

where $f(t) = jp(t) + d(t)$ is a complex function, $p(t)$ is the resonance peak $p(t)$, and $d(t)$ is a dumping coefficient. The fitting of one or more of these functions may be discretized in time in a number of parameters p^i, d^i corresponding to fitting times t^i .

According to some implementations, R may be assumed to be $R(t) = 1 - j\omega(t)$, which corresponds to a high pass filter.

The Fourier transform of the glottal pulse G may remain fairly constant over time. In some implementations, $G = g(t)g^*(t)$. The frequency profile of G may be approximated in a nonparametric fashion by interpolating across the harmonics frequencies at different times.

Given the model for the harmonic amplitudes set forth in EQN. 9, the model parameters may be regressed using the sum of squares rule as:

$$[A(t), \hat{g}(t), f_r(t)] = \underset{A(t), g(t), f_r(t)}{\text{argmin}} \left(A_h(t) - A(t)G(g(t), \omega(t)) \left[\prod_{r=1}^{N_f} F(f_r(t), \omega(t)) \right] R(\omega(t)) \right) \Big|_{\omega(t)=\hat{\phi}(t)h}^2 \quad \text{EQN. 11}$$

The regression in EQN. 11 may be performed in a nonlinear fashion assuming that the various time dependent functions can be interpolated from a number of discrete points in time. Because the regression in EQN. 11 depends on the estimated pitch, and in turn the estimated pitch depends on the harmonic amplitudes (see, e.g., EQN. 8), it may be possible to iterate between EQN. 11 and EQN. 8 to refine the fit.

In some implementations, the fit of the model parameters may be performed on harmonic amplitudes only, disregarding the phases during the fit. This may make the parameter fitting less sensitive to the phase variation of the real signal and/or the model, and may stabilize the fit. According to one implementation, for example:

$$[A(t), \hat{g}(t), f_r(t)] = \underset{A(t), g(t), f_r(t)}{\text{argmin}} \left(\|A_h(t)\| - \left\| A(t)G(g(t), \omega(t)) \left[\prod_{r=1}^{N_f} F(f_r(t), \omega(t)) \right] R(\omega(t)) \right\| \right) \Big|_{\omega(t)=\hat{\phi}(t)h}^2 \quad \text{EQN. 12}$$

In accordance with some implementations, the formant estimation may occur according to:

9

$$[A(t), f_r(t)] = \quad \text{EQN. 13}$$

$$\underset{A(t), f_r(t)}{\operatorname{argmin}} \left(\sum_h \operatorname{Var}_t \left(\frac{A_h(t)}{A(t) \left[\prod_{r=1}^{N_f} F(f_r(t), \omega(t)) \right] \right) \right)_{\omega(t)=\frac{d\Phi(t)}{dt}h} \right)^2.$$

EQN. 10 may be extended to include the pitch in one single minimization as:

$$[\Phi(t), A(t), f_r(t)] = \quad \text{EQN. 14}$$

$$\underset{\Phi(t), A(t), f_r(t)}{\operatorname{argmin}} \left(\sum_h \operatorname{Var}_t \left(\frac{s(t) \setminus M(\Phi(t))}{A(t) \left[\prod_{r=1}^{N_f} F(f_r(t), \omega(t)) \right] \right) \right)_{\omega(t)=\frac{d\Phi(t)}{dt}h} \right)^2.$$

The minimization may occur on a discretized version of the time-dependent parameter, assuming interpolation among the different time samples of each of them.

The final residual of the fit on the HAM($A_h(t)$) for both EQN. 10 and EQN. 11 may be assumed to be the glottal pulse. The glottal pulse may be subject to smoothing (or assumed constant) by taking an average:

$$G(\omega) = E_t(G(\omega, t)) = E_t \left(\frac{A_h(t)}{A(t) \left[\prod_{r=1}^{N_f} F(f_r(t), \omega) \right] \right)_{\omega=\frac{d\Phi(t)}{dt}h} \right). \quad \text{EQN. 15}$$

The reconstruction module **110** may be configured to reconstruct the speech component of input signal **114** with the noise component of input signal **114** being suppressed. The reconstruction may be performed once each of the parameters of the formant model has been determined. The reconstruction may be performed by interpolating all the time-dependent parameters and then resynthesizing the waveform of the speech component of input signal **114** according to:

$$\hat{s}(t) = 2\Re \quad \text{EQN. 16}$$

$$\left(\sum_{h=1}^{N_h} A(t)G(\omega) \left[\prod_{r=1}^{N_f} F(f_r(t), \omega(t)) \right] R(\omega(t)) \right)_{\omega(t)=\frac{d\Phi(t)}{dt}h} e^{j\Phi(t)}.$$

The output module **112** may be configured to transmit an output signal **118** to a destination **120**. The output signal **118** may include the reconstructed speech component of input signal **114**, as determined by EQN. 13. The destination **120** may include a speaker (i.e., an electric-to-acoustic transducer), a remote device, and/or other destination for output signal **118**. By way of non-limiting illustration, where communications platform **102** is a mobile communications device, a speaker integrated in the mobile communications device may provide output signal **118** by converting output signal **118** to sound to be heard by a user. As another illustration, output signal **118** may be provided from communications platform **102** to a remote device. The remote device may have its own speaker that converts output signal **118** to sound to be heard by a user of the remote device.

In some implementations, one or more components of system **100** may be operatively linked via one or more elec-

10

tronic communication links. For example, such electronic communication links may be established, at least in part, via a network such as the Internet, a telecommunications network, and/or other networks. It will be appreciated that this is not intended to be limiting, and that the scope of this disclosure includes implementations in which one or more components of system **100** may be operatively linked via some other communication media.

The communications platform **102** may include electronic storage **122**, one or more processors **124**, and/or other components. The communications platform **102** may include communication lines, or ports to enable the exchange of information with a network and/or other platforms. Illustration of communications platform **102** in FIG. **1** is not intended to be limiting. The communications platform **102** may include a plurality of hardware, software, and/or firmware components operating together to provide the functionality attributed herein to communications platform **102**. For example, communications platform **102** may be implemented by two or more communications platforms operating together as communications platform **102**. By way of non-limiting example, communications platform **102** may include one or more of a server, desktop computer, a laptop computer, a handheld computer, a NetBook, a Smartphone, a cellular phone, a telephony headset, a gaming console, and/or other communications platforms.

The electronic storage **122** may comprise electronic storage media that electronically stores information. The electronic storage media of electronic storage **122** may include one or both of system storage that is provided integrally (i.e., substantially non-removable) with communications platform **102** and/or removable storage that is removably connectable to communications platform **102** via, for example, a port (e.g., a USB port, a firewire port, etc.) or a drive (e.g., a disk drive, etc.). The electronic storage **122** may include one or more of optically readable storage media (e.g., optical disks, etc.), magnetically readable storage media (e.g., magnetic tape, magnetic hard drive, floppy drive, etc.), electrical charge-based storage media (e.g., EEPROM, RAM, etc.), solid-state storage media (e.g., flash drive, etc.), and/or other electronically readable storage media. The electronic storage **122** may include one or more virtual storage resources (e.g., cloud storage, a virtual private network, and/or other virtual storage resources). The electronic storage **122** may store software algorithms, information determined by processor(s) **124**, information received from a remote device, information received from source **116**, information to be transmitted to destination **120**, and/or other information that enables communications platform **102** to function as described herein.

The processor(s) **124** may be configured to provide information processing capabilities in communications platform **102**. As such, processor(s) **124** may include one or more of a digital processor, an analog processor, a digital circuit designed to process information, an analog circuit designed to process information, a state machine, and/or other mechanisms for electronically processing information. Although processor(s) **124** is shown in FIG. **1** as a single entity, this is for illustrative purposes only. In some implementations, processor(s) **124** may include a plurality of processing units. These processing units may be physically located within the same device, or processor(s) **124** may represent processing functionality of a plurality of devices operating in coordination. The processor(s) **124** may be configured to execute modules **104**, **108A**, **108B**, **110**, **112**, and/or other modules. The processor(s) **124** may be configured to execute modules **104**, **108A**, **108B**, **110**, **112**, and/or other modules by software; hardware; firmware; some combination of software,

11

hardware, and/or firmware; and/or other mechanisms for configuring processing capabilities on processor(s) **124**.

It should be appreciated that although modules **104**, **108A**, **108B**, **110**, and **112** are illustrated in FIG. **1** as being co-located within a single processing unit, in implementations in which processor(s) **124** includes multiple processing units, one or more of modules **104**, **108A**, **108B**, **110**, and/or **112** may be located remotely from the other modules. The description of the functionality provided by the different modules **104**, **108A**, **108B**, **110**, and/or **112** described below is for illustrative purposes, and is not intended to be limiting, as any of modules **104**, **108A**, **108B**, **110**, and/or **112** may provide more or less functionality than is described. For example, one or more of modules **104**, **108A**, **108B**, **110**, and/or **112** may be eliminated, and some or all of its functionality may be provided by other ones of modules **104**, **108A**, **108B**, **110**, and/or **112**. As another example, processor(s) **124** may be configured to execute one or more additional modules that may perform some or all of the functionality attributed below to one of modules **104**, **108A**, **108B**, **110**, and/or **112**.

FIG. **3** illustrates a method **300** for performing voice enhancement and/or speech features extraction on noisy audio signals, in accordance with one or more implementations. The operations of method **300** presented below are intended to be illustrative. In some embodiments, method **300** may be accomplished with one or more additional operations not described, and/or without one or more of the operations discussed. Additionally, the order in which the operations of method **300** are illustrated in FIG. **3** and described below is not intended to be limiting.

In some embodiments, method **300** may be implemented in one or more processing devices (e.g., a digital processor, an analog processor, a digital circuit designed to process information, an analog circuit designed to process information, a state machine, and/or other mechanisms for electronically processing information). The one or more processing devices may include one or more devices executing some or all of the operations of method **300** in response to instructions stored electronically on an electronic storage medium. The one or more processing devices may include one or more devices configured through hardware, firmware, and/or software to be specifically designed for execution of one or more of the operations of method **300**.

At an operation **302**, an input signal may be segmented into discrete successive time windows. The input signal may convey audio comprising a speech component superimposed on a noise component. The time windows may include a first time window spanning a duration greater than a sampling interval of the input signal. Operation **302** may be performed by one or more processors configured to execute a preprocessing module that is the same as or similar to preprocessing module **106**, according to some implementations.

At an operation **304**, a transform may be performed on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows. A first sound model may describe a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal. Pitch may be the rate of change of phase over time. Chirp may be the rate of change of pitch over time. Operation **304** may be performed by one or more processors configured to execute a transform module that is the same as or similar to transform module **108A**, according to some implementations.

At an operation **306**, linear fits in time of the sound models over individual time windows of the input signal may be obtained. The linear fits may a first linear fit in time of the first

12

sound model over the first time window. Operation **306** may be performed by one or more processors configured to execute a transform module that is the same as or similar to transform module **108A**, according to some implementations.

Although the present technology has been described in detail for the purpose of illustration based on what is currently considered to be the most practical and preferred implementations, it is to be understood that such detail is solely for that purpose and that the technology is not limited to the disclosed implementations, but, on the contrary, is intended to cover modifications and equivalent arrangements that are within the spirit and scope of the appended claims. For example, it is to be understood that the present technology contemplates that, to the extent possible, one or more features of any implementation can be combined with one or more features of any other implementation.

What is claimed is:

1. A system configured to perform voice enhancement and/or speech features extraction on noisy audio signals, the system comprising:

a memory storing computer executable instructions; and one or more processors coupled to the memory and configured to execute the computer executable instructions to:

segment an input signal into discrete successive time windows, the input signal conveying audio comprising a speech component superimposed on a noise component, the time windows including a first time window spanning a duration greater than a sampling interval of the input signal;

perform a transform on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows, the sound models including a first sound model including a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal, pitch being the rate of change of phase over time, chirp being the rate of change of pitch over time; and obtain linear fits in time of the sound models over individual time windows of the input signal, the linear fits including a first linear fit in time of the first sound model over the first time window.

2. The system of claim 1, wherein a linear regression is used to fit the first sound model over the first time window to obtain the first linear fit.

3. The system of claim 1, wherein the first model is a superposition of harmonics in the first time window with a linearly varying fundamental frequency.

4. The system of claim 1, wherein the one or more processors are further configured to execute the computer executable instructions to impose continuity in a pitch estimation of the first sound model.

5. The system of claim 1, wherein harmonic amplitudes in the first sound model are piecewise linear and/or continuous in time.

6. The system of claim 1, wherein an integral phase of the first sound model is optimized via a nonlinear regression.

7. The system of claim 1, wherein the integral phase is optimized via multiple iterations of the nonlinear regression.

8. The system of claim 1, wherein a regression to estimate the integral phase is performed locally.

9. The system of claim 1, wherein the integral phase is approximated with a number of time points to reduce the degrees of freedom.

13

10. A processor-implemented method to perform voice enhancement and/or speech features extraction on noisy audio signals, the method comprising:

segmenting, using one or more processors, an input signal into discrete successive time windows, the input signal conveying audio comprising a speech component superimposed on a noise component, the time windows including a first time window spanning a duration greater than a sampling interval of the input signal;

performing, using one or more processors, a transform on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows, the sound models including a first sound model including a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal, pitch being the rate of change of phase over time, chirp being the rate of change of pitch over time; and

obtaining, using one or more processors, linear fits in time of the sound models over individual time windows of the input signal, the linear fits including a first linear fit in time of the first sound model over the first time window.

11. The method of claim 10, wherein a linear regression is used to fit the first sound model over the first time window to obtain the first linear fit.

12. The method of claim 10, wherein the first model is a superposition of harmonics in the first time window with a linearly varying fundamental frequency.

13. The method of claim 10, further comprising imposing continuity in a pitch estimation of the first sound model.

14. The method of claim 10, wherein harmonic amplitudes in the first sound model are piecewise linear in time.

15. The method of claim 10, wherein an integral phase of the first sound model is optimized via a nonlinear regression.

14

16. The method of claim 10, wherein the integral phase is optimized via multiple iterations of the nonlinear regression.

17. The method of claim 10, wherein a regression to estimate the integral phase is performed locally.

18. The method of claim 10, wherein the integral phase is approximated with a number of time points to reduce the degrees of freedom.

19. One or more non-transitory computer readable storage media encoded with instructions that, when executed by a processor, cause the processor to:

segment an input signal into discrete successive time windows, the input signal conveying audio comprising a speech component superimposed on a noise component, the time windows including a first time window spanning a duration greater than a sampling interval of the input signal;

perform a transform on individual time windows of the input signal to obtain corresponding sound models of the input signal in the individual time windows, the sound models including a first sound model including a superposition of harmonics sharing a common pitch and chirp in the first time window of the input signal, pitch being the rate of change of phase over time, chirp being the rate of change of pitch over time; and

obtain linear fits in time of the sound models over individual time windows of the input signal, the linear fits including a first linear fit in time of the first sound model over the first time window.

20. The non-transitory computer readable storage media of claim 19, wherein an integral phase of the first sound model is optimized via a nonlinear regression.

* * * * *